UNIVERSITY OF QUEENSLAND
ST. LUCIA
QUEENSLAND
AUSTRALIA 4067

PRENTICE
COMPUTER
CENTRE

NEWSLETTER

N-242

17-Aug-79

CURRENT NETWORK DEVELOPMENTS

Authorised by the Director of the Computer Centre

N-242
17-Aug-79

Networks - Director's Comments

In May 1978 a special newsletter was issued on network developments. Mr. Barker, Systems Analyst of the Centre, has prepared a further paper so that users may be aware of current facilities and progress. We welcome comments and advice.

There are a number of separate areas of development.

Currently the transmission system uses TELECOM Australia telephone lines installed on campus. This means that we *must* hire from TELECOM, the modems at each end of the transmission link. The cost of these modems increases with the speed of transmission. Additionally TELECOM Australia places restrictions on service offerings which inhibit appropriate developments within the University. Having regard to the present level of data communications use on campus and the expected developments in the future, the Centre felt that a case could exist to replace the TELECOM Australia data transmission facilities with University facilities. A preliminary study indicated that there would be economic and service advantages and a report has been referred to the Director Buildings and Grounds. Our hope is that this development will be approved in principle following which it would seem appropriate to convene a technical committee with representatives from Electrical Engineering, Department of Computer Science, Buildings and Grounds and the Computer Centre. The aim would be to advance further with the technical design of the transmission system to enable more definite specifications and costs. Another area of development relates to the use of the transmission system for the various applications required to enable the effective distribution and association of computing intelligence throughout the network. There is now a considerable body of knowledge and experience arising out of network developments in North America and Europe and appropriate international standards have been agreed at a surprising speed. Unfortunately we do not have the resources to develop the type of network we would like from the outset. The network software available from DEC has provided some problems as there is a lack of protocol homogeniety as between PDP11 network software and PDP10 network software. Nevertheless by co-operation with other DEC10 users overseas and Australia the facilities available to users should soon improve considerably. These are described in Mr. Barker's paper.

The number of terminals with increasing levels of intelligence will grow and there will be increasing traffic on the network for file transfers between systems and down line loading of programs. We have provided facilities whereby terminals connected directly to the communications front-end of the KL10 can connect to the KA system. Shortly we will release a facility to enable PDP11/34 RSX terminal users to have virtual terminal facilities to the KL10. This system also provides that output from the KL10 standard line printer spooler can be returned to the 11/34 without interfering with the local line printer spooling of the RSX11-M operating system. In the latter part of the year we will be testing our network node software. This will allow the user of any terminal connected to the node to be able to set host to any host computer on the network where the user has arranged privileges. There are of course many areas of development and improvement (eg. network management so that users and the Centre know what is going on in the network, incorporation of PDP11 Unix Operating Systems etc.), but our pace is necessarily limited by resources available.

Another major area of development is the ability to connect to other networks. Towards the end of this year we will be testing a gateway node developed by CSIRO which will permit the University Network to be connected to the CSIRO network. There will no doubt be many problems particularly in relation to network management and charging. Apart from giving users access to software and other facilities not available on the University system, it heralds the start of a national network for University and scientific computers.

Additionally, we have made arrangements for connection of our system to the Overseas Telecommunications Commission MIDAS system which will provide access into the U.S. Tymet and Telenet networks. In the long run perhaps the greatest advantage to the University from the linking to other networks will be the potential for co-operative research nationally and internationally.

My hope is that Mr. Barker's paper provides a basis for Departments to consider their data communications needs and discuss their requirements with us. The Computing Policy Committee has provided some funds to cover the capital cost of interface equipment to allow the linking of mini-computers to the central computing facility. Please contact Mr. Barker, Mr. Hartwig or myself.

Alan W. Coulter
extension 2189

Introduction

In May of 1978 we devoted a complete issue of the Newsletter to the subject of the development of a network to link the growing number of mini and microcomputers and the central DEC PDP10. We explored various likely user requirements and outlined the then current state of knowledge about the hardware and software components that might be needed to meet such requirements. Finally we asked readers to comment and make suggestions about our proposals.

Now, some sixteen months later we are reviewing the subject again. During the ensuing period the number of installed minicomputers has grown to over forty and we have acquired a better understanding of their relative strengths and we have a better idea of their proper areas of application. We have acquired some useful network software from various sources and modified it where necessary to suit our requirements. We are developing some software ourselves.

DEC have upgraded their standard network software and reduced some of the incompatibilities between their own two network systems - DECnet-11 and ANF-10, which we discussed in the previous newsletter.

In an endeavour to reduce recurrent costs of computing, at least on campus, the Centre has proposed some alternative means for linking terminals and other devices to the network other than via TELECOM Australia facilities. The estimated future lease costs for on-campus users appears to be sufficiently high to warrant some capital investment in appropriate alternatives.

Finally as a measure of the importance which they attach to the network philosophy, the Computing Policy Committee has made funds available to be used to assist users to acquire the necessary hardware to connect to the network. These funds are to provide for both the equipments required at user sites and and equipment required at the central site to complete the connection.

In this issue, in response to requests from readers of our previous Newsletter, we are including some more basic explanatory sections on network operations.

1.0 WHY A NETWORK

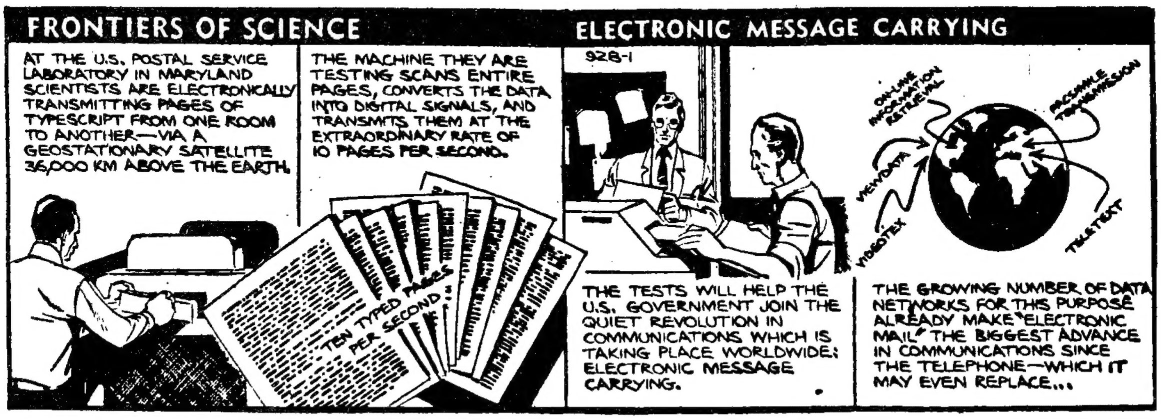

Computer users in a network enjoy flexibility and choice of facilities to suit varying processing requirements. Computers installed on campus range from microprocessors through a range of minicomputers to the central KL10. The microcomputers are usually dedicated to tasks such as experiment monitoring and data logging, but the processing of captured data generally requires access to a computer of greater capacity. The minicomputers are capable of performing many preparatory functions such as editing of data files, program development, data reductions and other not too ambitious processing, but they cannot offer the range of software and processing power of the KL10. In the absence of suitable linking mechanisms the transfer of data between different computers is often difficult and frustrating and involves expenditure on alternative transfer media such as disk or magnetic tape equipment anyway. Access to different computers via a user's terminal can be arranged by switching the terminal between lines connected to the respective computers but without the ability to transfer data between the computers, this facility is of limited use.

Another possibility with connected computers is to set up co-operating processes so that a program on one machine controls a function in another machine. For example, if there is limited storage in a microcomputer engaged in data logging then the data may be transferred as it is collected to another minicomputer which has appropriate capacity. Or, an experiment being monitored by a microcomputer can be controlled in real time, by a program running in a larger minicomputer or the KL10, which is processing data being collected and transmitted from the microprocessor. These are examples of a task-to-task communications facility provided within the network.

An important related aspect is the inevitable loading to saturation of the KL10, a situation which users of the KA10 will remember from the period just prior to the installation of the KL10. An obvious way to alleviate the problem is to encourage as much editing and other preprocessing of files on minicomputers and intelligent terminals as is possible, thereby freeing up the KL10's resources for more demanding computing. Microprocessor based intelligent terminals with some

local storage are becoming cheaper and may be connected to the networks in the future, Access to central file storage media also reduces the need for local disk capacity. Connected users have access to public disk areas or to private disk packs or to other data bases on the KL10.

A further likely development, as a consequence of the decreasing cost of computer hardware, will be the establishment of computers dedicated to performing single, but generally used tasks, to provide facilities for all network users.

Beyond the University's network there are other similar networks, for example CSIRONET. CSIRO and the Centre are currently developing, as a joint project, the necessary means of connecting both networks. This will mean that users connected to either network will be able to access both network's facilities. The project is scheduled for completion before the end of this year.

Another recently announced service should be of interest to network users. The Australian Overseas Telecommunications Commission (OTC) has installed a computer which is a node, via satellite and overseas cable, in the TYMNET and TELENET networks in the U.S.A. These networks include over 200 host computers offering a wide range of services and data bases to time sharing users. Any person who has terminal access to the OTC node via TELECOM STD links can gain immediate access to any of the overseas hosts with whom they have made prior arrangements. This service is known as MIDAS (*M*utimode *I*nternationnal *D*ata *A*cquisition *S*ervice).

Normally, users would need an acoustic coupled terminal to make such a connection, but the Centre is developing a mechanism whereby any of our users can connect directly to MIDAS through the University's network. There are potential cost savings also from volume discounts where the Centre can be billed centrally on behalf of several otherwise separate users.

Electronic mailing is another useful service which can be implemented via an existing network. Such a service allows anyone connected to the network to send messages to any or all other connected users. Messages include interdepartmental memos, letters or any other information currently typed and distributed by internal mail despatch. This is a store and forward system. Users enter messages into their terminals and the network stores the messages until someone at the destination activates his terminal, at which point he is presented with a list of waiting messages. He may display them on his screen or direct them to a hard copy printer terminal for later reading. Electronic mail is very fast and efficient, offering savings in time, and paper and despatching costs.

2.0 NETWORK BASICS

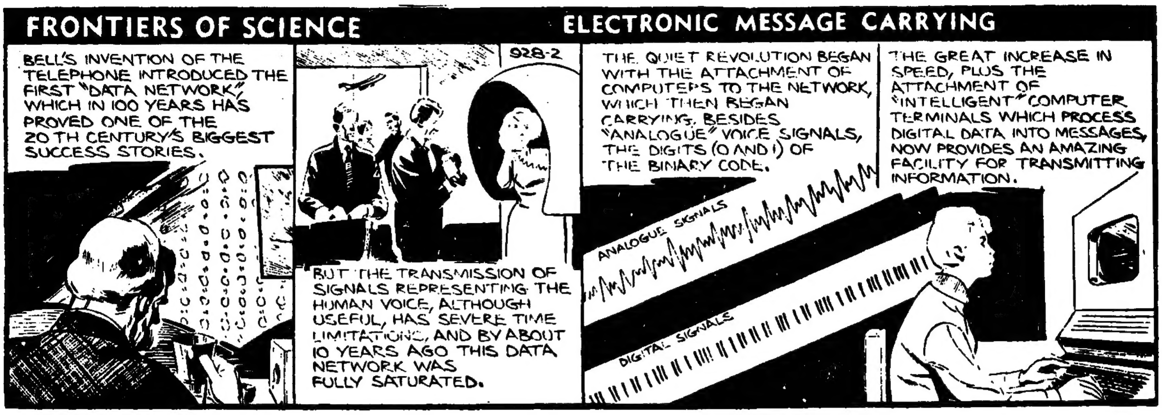

The diagram shows a network similar to that envisaged for the University. The network's structure is not unlike the telephone system in that there are a number of interconnected 'nodes' which are like local exchanges to which all local terminals and other input or output devices such as line printers are connected. The nodes in turn are connected by main trunks which carry data at high speed in a concentrated format. Where nodes are geographically dispersed they are connected via leased TELECOM lines; where nodes are adjacent to one another, for example in the one building, they are usually directly connected via private lines.

Some of the nodes perform computing functions, ie. they have operating systems and do computing work as users require. for example, the KL10 and departmental minicomputers. Computing nodes are commonly referred to as 'hosts'. Other nodes may also consist of minicomputers but their function is solely to handle data traffic in the networks. These nodes are set up separately in order to free the computing nodes of the considerable load involved in communications control.

There are different types of traffic on the network, for example terminal traffic between users and computers or between terminals, data files being transferred between computers, data being output to remote printers or input from remote card readers and information about the state of the network itself.

At present, in order to connect to the network, every user must lease a TELECOM line and a pair of special terminating devices called 'modems'. Modems are devices which change digital signals to voice frequency signals for transmission of data over the telephone network and also isolate the telephone system from unwanted voltages or signals from non-standard equipment. On

campus there is no lease charge for telephone circuits which are owned by the University, but there is still a modem lease charge. Off campus there is always a line lease charge.

Because all communications links, especially TELECOM circuits, are subject to various noise conditions which cause data errors, it is necessary to employ special error detection and correction methods to ensure that transmitted data is correctly received at its destination. Also, methods are required to ensure correct sorting out of data where trunks are shared by several nodes. These methods are embodied in special communications hardware and or software and the set of rules defining exactly how data is transferred and errors are detected and corrected in a particular network is known as the communications 'protocol'.

When a user at a terminal types a command it is echoed by his local node and typed back on the terminal. He can see immediately whether or not the command was transmitted successfully. If it was not, he types it again. This simple form of error detection and correction is fine for users of terminals. However, for transmission of data between computers it is quite inadequate. For example when transferring files or directing output to a remote printer the transmission speed is such that error detection and retransmission must be automatic and efficient.

The method usually employed is to break up data to be transmitted into small parcels of a few hundred or so characters and to add to each parcel a group of control characters, and to transmit each parcel as a separate message. The control characters are said to form an 'envelope' around the data. The control characters in each message contain information necessary for the communications system to guarantee delivery of each message correctly. This information includes codes for the source and destination, the message number, checksums on the enclosed data, and so on.

As each message is sent, the destination node transmits confirmation of correctness and receipt, and if confirmation does not arrive then the source retransmits the message. Eventually, the destination node can reassemble the original data from the large number of small messages. For example, if a user directs output to a remote printer, then each line of print may be sent as a separate message wrapped in an envelope of control information.

The format and meaning of the control characters and the procedures for the control of message flow are determined by the particular protocol in use. In practice users should not have to worry about communications control; it is normally handled transparently by the systems software.

In addition to the modem, another piece of equipment is usually required to act as an interface between the modem and the users equipment. This is known as a communications interface or controller or line driver. It is actually a part of the users equipment but it must be acceptable to TELECOM for connection to a modem. Generally, all terminals have interfaces inbuilt, but for other computing equipment such as minicomputers and microprocessors, it will generally be necessary that they be purchased separately. The choice of interface is determined by the mode and speed of transmission required.

There are two primary modes of transmission; asynchronous and synchronous. In asynchronous transmission, data is transmitted one character at a time at varying intervals with additional start and stop bits for control of timing. In synchronous transmission, data is transmitted in blocks of characters at a fixed rate with the transmitter and receiver synchronized, thus eliminating the need for extra start and stop bits.

Asynchronous transmission is simpler and cheaper to implement but it is not very efficient. It is used, for instance, between terminals and nodes and between nodes with light traffic. Synchronous transmission is more efficient but also more expensive to implement. It is used for all major trunk links in the network.

Obviously, there is a good deal of overhead involved in error free transmission in terms of non-data characters transmitted and computing overhead. In order to offset the overheads and to move data in practical acceptable times, speed of transmission should be as high as possible. Typical speed ranges available currently in our network are 300 to 1200 bits per second for asynchronous transmission and 1200 bits up to 9600 bits per second for synchronous transmission. The principal determinant of choice of speed however is cost. The costs of modem and line rental and communications interfaces all rise significantly with speed. Table I contains current TELECOM charges, Table II contains current costs of some available communications interfaces.

Because individual high speed lines are expensive various methods for sharing lines have been developed. Terminal concentrators are nodes in a network, incorporating a micro or minicomputer, whose sole function is to funnel data to and from a number of local terminals, down a single high speed line. The technique used for simultaneous transmission of data from several sources on a shared line is called multiplexing. On the other end of the line there must of course be some form of demultiplexing. A remote batch station, or remote job entry (RJE) station, is a form of multiplexer where a remote printer and card reader are connected via a minicomputer to a single high speed line. Remote batch stations may also incorporate terminal multiplexing. Users of a remote batch station or terminal multiplexer, should enjoy the same facilities as, and with only slightly slower response, than users connected directly to the host.

Nodes in the network usually consist of either host computers running an operating system incorporating a command decoder, for example TOPS-10 or RSX-11, or of mini or micro computers acting as some sort of concentrator. In the latter case their existence in the network is transparent to users since they cannot respond to user commands.

A useful additional function is provided by adding a few very limited commands to a concentrator to enable it to act as a network switch. Using these commands a user may direct that he be connected virtually to any host in the network or to enquire about the state of the network in general. For example, he may wish to know which particular hosts are active at any particular time. Switching nodes simply examine destination addresses in message headers and pass the messages on along the appropriate path.

The performance and reliability of the network as a whole is as important as that of any single node in it. In a network comprising various communications nodes and multiple data paths, the problem of network maintenance is complex. There is a need to continuously monitor traffic flows to detect any breakdown of data paths or unacceptable error rates and congestion. When problems do arise, it must be possible to run diagnostics to determine the cause. In any case, statistics on network use will be required for management and accounting purposes. In even a moderately complex network, the magnitude of the task will usually require a special node dedicated to network management.

In respect to the performance of the network overall a major concern is the ability of the communications system to handle all the overhead involved in multiplexing, routing, switching etc., in addition to simply moving data about. The final test is the response of the system to a user sitting at a terminal trying to do useful work. A critical factor in response time is the speed of the interconnecting telephone lines and this is determined by modem speeds in respect of which, in turn, cost is the limiting factor. It is important therefore that at least on campus, the University does proceed with the laying of private ducts and data cables which will allow implementation of economic high speed links.

3.0 NETWORK REQUIREMENTS

In this section we discuss what we believe to be the facilities which are necessary for users to make good use of the network.

Any terminal user should be able to effectively attach to any host in the network. For example, a user whose terminal is connected physically to a departmental minicomputer should be able to obtain a connection via the network to the KL10 or to CSIRONET or to another minicomputer. He can then log onto whatever computer is appropriate for the work he wants to do. This is the 'virtual terminal' facility. This facility will become more important with the addition in the future of computers performing specialized functions.

Security remains an important issue. Users should enjoy the same level of security that is currently available to KL10 or KA10 users. In the network environment this means that users connected to one node should not be able to connect to other nodes or transfer data to or from other nodes wherein they have not established accounts, passwords etc., or have not made some other arrangement regarding file protection with a user on the destination node.

The ability to transfer data between computers is most important. It is usually performed by system software and is accomplished on behalf of the user by co-operating tasks in the source and destination hosts. Security considerations should require that the user be logged in on one host and to have to log in on the other before being able to initiate a transfer. Once logged-in then normal access privileges apply. Logging in on the remote hosts will require a virtual terminal facility. Alternatively, rather than log in to the second host, the user may have to quote account details, password, etc., as part of the request command line. Transfers of data need not happen immediately, requests may be queued. Queuing of requests has the benefit of enabling the system to cope with peaks in network

traffic and a priority schedule would allow transfers to be deferred till a lower cost time of day, for example, overnight.

The submission of command files is a variation on data file transfers. For example, the creation of a control file for the KL10 on a minicomputer and transferring it to the KL10 for execution. This facility requires an additional system task in the destination node to recognize the arrival of a command file and to submit it to the normal batch queue. The same security constraints apply as for data file transfers.

Access to remote peripheral devices is useful, especially for users who are doing work on a remote host and who want output produced on a local printer or terminal. For users connected to a remote batch station-concentrator, the facility should be standard - that is, part of the batch station function. For users connected to a minicomputer running under its own operating system, the facility requires the installation of a special spooler in the host node and a task in the local node to control the output device.

Finally, information about the general state of the network is necessary from time to time, both to users and to those who have to administer the network. Such information includes names and status of all nodes currently connected.

4.0 THE CURRENT NETWORK

The diagram shows the network as it now exists and as it is likely to develop. The network is drawn as a series of nodes (the boxes) connected by TELECOM or private lines. The different types of node are indicated by the codes inside the boxes. Individual nodes are known to the network software by unique names and numbers which are shown thus - *UQKL10*(1). The various network commands may use either the node name or the number.

The largest nodes are the KL10 and the KA10 host computers. All access to the KL10 is via one or other of the two DN87 nodes which are so called 'front-end' communications processors. Their function is to process all terminal and network traffic destined for the KL10 and to reroute traffic destined for other nodes. They each contain a PDP11/40 processor. Most terminal users are currently connected directly to a DN87.

The nodes labelled DN82 are terminal concentrator-remote batch stations. There is one installed in the Commerce building - *COMERC*(30), and at Griffith University - *GUAES*(50). Users connected to a DN82 have the same facilities as those connected directly to a DN87. If they need access to another node, for example the KA10, then they will be switched through a DN87. A DN82 does not have a command decoder, all commands issued by a connected user are executed by the host to which they are currently connected. A DN82 node can also select the best path to the destination node.

The DN82 node named *SSPRIN*(20) has the special function of running the low cost self service printer located in the Hawken Batch Station. This station allows users of either the KA10 or KL10 to queue output (which is less than 20 pages long) to that printer and help themselves to that output.

The DN83 nodes have the same facilities as a DN82 but in addition they have extra code to enable connection to 'foreign' hosts. For example, the node labelled *CSIRO*(21). This node looks like a standard DN82 node to the University network but appears to CSIRONET as a standard CSIRONET node. The extra code translates between the two network protocols.

There are no DN83 nodes installed currently. Their principal function is to act as network switchers when traffic in the network builds to a level which causes undue loading on the DN87's. Lines to external nodes, such as the CSIRO node, may be connected to a DN87 or a DN83 depending on loading factors.

Nodes labelled 'U-mux' and 'PATX' (*P*rivate *A*utomatic *T*erminal *EX*change) are low cost microprocessor based terminal concentrators or switches.

U-mux is a simple concentrator to allow several (up to about eight) terminals to share a common line to improve circuit utilization. They perform multiplexing with simple communications protocol. They are particularly cost effective if the shared line is a TELECOM line requiring leased modems. The shared line is terminated in any DN82 or DN83 series node where demultiplexing takes place. The service to connected users is the same as if they were connected directly to the node.

PATX is an extension of U-mux in that in addition to multiplexing to a shared line, the user may be switched if he requires to an alternative output line. These devices find application where there are terminal users who need an inexpensive facility to switch between local minicomputer and the main network. For example, a terminal laboratory of six terminals which needs to be connected to a local PDP-11 for some assignments and to the KL10 for others.

An extension of the idea would allow a building to be fitted with say twenty four terminal outlets in various parts of the building and into which could be plugged up to eight terminals at any one time. This would provide a means of setting up terminal laboratories in a fairly flexible way with the minimum number of terminals.

PATX and U-mux are not standard DEC products. Devices to perform PATX type functions are available but they are overpriced and do not fit neatly into our network scheme. The Computer Centre is currently designing a prototype microprocessor module which will become the basis of the concentrators and switches which will exactly meet our requirements as the network develops.

The remaining nodes in the diagram represent the various PDP-11 minicomputers in departments. They usually run RSX-11M or RT-11 in 11/34's or 11/03's. Some possible interconnections are shown.

Protocols

A network comprising PDP-10's and the other DN80 series nodes described, uses standard DEC hardware and DEC network software known as ANF-10. All links between these nodes use the communications protocol defined within ANF-10.

A network comprising PDP-11's uses another DEC standard network known as DECnet-11. Regrettably it is incompatible with ANF-10 in certain important areas. This is regrettable because a very natural network configuration is one incorporating PDP-10's and PDP-11's - it is one which offers a nice hierarchy of facilities.

DEC communications protocols are multilayered. The lowest layer is concerned with maintenance of control over a physical link between two points. For example, between any two nodes. This physical link control uses the Digital Data Communications Message Protocol (DDCMP) which defines the format and meaning of the message header and the procedures for control of message flow.

The second layer is concerned with logical links, several of which may operate simultaneously over the one physical link. This layer allows several users of any node to share a common link with any other node and also controls routing of messages through the network. It is implemented by another envelope of control characters.

ANF-10 and DECnet-11 both use DDCMP, but they differ in their logical link protocols. ANF-10 uses a protocol known as Network Command Language (NCL) and DECnet-11 uses Network Service Protocol (NSP).

To overcome the problem, DEC have provided a software module called the DECnet Compatible Port (DCP) which is designed to reside in any DN80 series node and to translate between NCL and NSP. It is only a partially successful solution and the incompatibility has caused difficulty in developing the necessary software to meet all network requirements. However, software development is proceeding and we are confident of overcoming most problems.

The software that is currently available will allow users to perform most immediately required tasks.

Available Software

A user connected to a PDP-11 in the network running RSX-11 or RT-11 and DECNET-11 has the following facilities as a standard :-

* He may send messages to other terminals on local or remote PDP-11 nodes using the terminal communications utility TLK;

* He may transfer data files between PDP-11 nodes, and transmit and submit command files to remote PDP-11 nodes using the Network File Transfer utility NFT;

* He may perform task-to-task communications between PDP-11 nodes using FORTRAN and MACRO-11 callable subroutines and the DECnet-11 intertask communications services, or between the PDP-11 and the KL-10 using the DECnet Compatible Port.

A user connected to a DN87 or other DN80 series node has the following ANF-10 facilities as standard :-

* He may obtain a virtual connection to the KA or KL10 or any other host connected via a DN83 to the network using the SET HOST command;

* He may send messages to other terminals on local or remote DN80 series nodes using the SEND command;

* He may gain access to devices connected to DN80 series remote nodes using the LOCATE command;

* He may perform task-to-task communications between PDP-10's, or between a PDP-10 and a PDP-11 using the DECnet Compatible Port;

* He may request information about the state of the network using the NODE and WHERE commands.

At the moment there is no available DECnet interface to the UNIX operating system, but it is the Centre's intention to provide an interface at the earliest opportunity.

For general communications between PDP-11's and DEC 10's there is little DEC supplied software. There are however a variety of user written programs designed to make use of the DECnet Compatible Port.

Immediately available is CLINK, a communications package developed by National Institutes of Health, USA, which provides a virtual terminal and file transfer capability to PDP-11 users connected via a DN80 series node to a DEC10.

Briefly, a user logs on to the PDP11 and runs a program on the 11 which connects him to the 10 as a virtual terminal. He can then log on to the 10, if the appropriate accounting entries have been set up, and perform all normal timesharing operations. In addition he can transfer files between the two systems by use of the CLINK program residing on the 10.

CLINK will support only one user at a time on an PDP-11 but is sharable on the 10. CLINK requires that the 11 terminal be the CTY or any other terminal connected via a DL11 interface and the line to the 10 requires a DL11 also. It is reliable and has been in use for some time.

The line speed should not be less than 1200 bits per second otherwise file transfer times are impractical. Thus, at least at the present time on campus, a pair of 1200 b/s modems are required also. CLINK does not use the DECnet software but DDCMP functions are performed in the program in the 10.

The Centre has acquired other software which performs similar functions but which uses standard network software. It is currently being modified to overcome some minor problems and will be available by the end of this year. This software will allow multiple virtual terminals for RSX users connected to DL11's or DZ11's on a PDP-11 connected in turn to any DN80 series node. Having obtained virtual connection to the KL10, they may log in and do any work including file transfers to or from their PDP-11. This new software does use the DEC network software and the result is a considerable saving in cost for file transfers because the protocol processing overhead is not carried by the user programs.

RSX users will also be able to spool output to a printer or terminal device connected to their PDP-11 whilst logged in on the KL10.

5.0 CONNECTING TO THE NETWORK

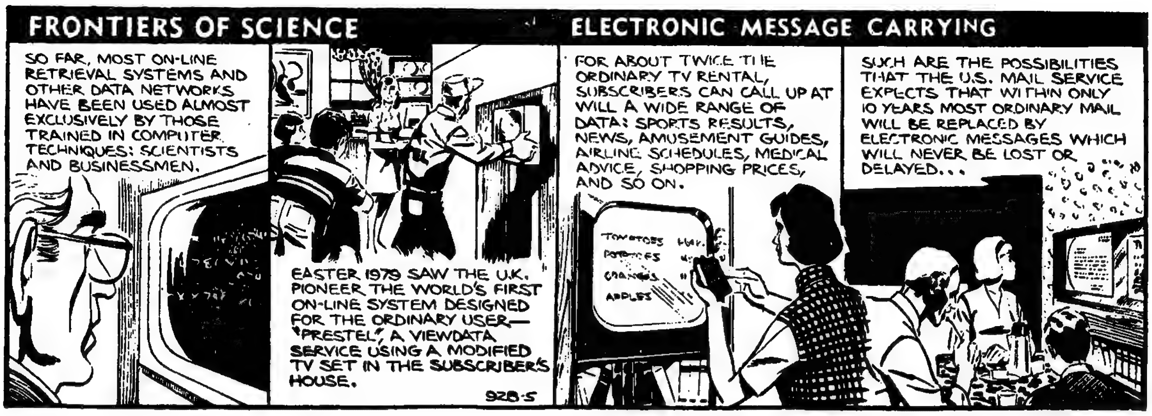

Anyone wishing to discuss possible connection to the network or any other related matter are invited to contact either John Barker (ext. 3016), or Arthur Hartwig (ext. 3021). We welcome any comment, query or suggestion. This section offers some preliminary guide to prospective network users.

Users wishing to connect to the network will need to select an appropriate communications interface and modems. Factors that should be considered include- estimates of data volumes, line speeds, CPU loading due to communications processing, upgrade path, costs.

There are three different interfaces currently available which between them, cover a reasonable range of price and speed, Their characteristics are summarised in Table II.

As a general guide, we suggest that a minimum line speed of 1200 b/s will be necessary for practical purposes. At this speed, the effective transmission rate will be around 90 characters per second; a 100 block file will take about 10 minutes to transmit.

CPU loading arises because of the processing involved with CRC error check calculations and general message synchronization. Some interfaces perform some or all of these functions, thus relieving the CPU of that load. For the three interfaces in Table II, estimated loadings for a PDP-11 are as follows (for an average message size of 256 characters):-

Interface	*CPU Loading*	*What is Done*
DL11	35%	all communications processing is performed in the PDP-11 CPU.
DUP-11	30%	CRC error checking is performed in the interface.
DMC-11	10%	complete DDCMP protocol performed in the interface.

The DL-11 is an asynchronous interface. If it is connected to a TELECOM line its maximum speed is 1200 b/s because that is the highest speed available in asynchronous modems. In order to upgrade to a higher speed line, a synchronous interface would be required. On the other hand, if private lines become available on campus then the DL-11 may be driven up to 9600 b/s since TELECOM modems would not be required.

The DUP-11 is a synchronous interface and the lowest speed synchronous modem available is 2400 b/s. It will be slightly more efficient than the DL-11 at the same speed, and it performs CRC-16 block check calculations and notifies the user program in the event of an error. Line speed may be upgraded to 9600 b/s.

The DMC-11 is a synchronous line interface which implements the complete DDCMP protocol in its inbuilt microprocessor. Line speed may be upgraded to 9600 b/s with a TELECOM modem, or up to 1,000,000 b/s with a private line.

Table I

On Campus Modem Installation and Rental Charges

Speed - Bits per second	Annual Rental	Installation	Synchronous/ Asynchronous
	$	$	
Up to 300	120	120	A
600/1200	420	144	A
2400	660	192	S
4800	780	216	S
9600	1440	300	S

These charges are 60% of normal off-campus charges because they do not include a charge for the line. That is, the charges are for the modem only because the University owns the telephone cables.

Off campus the rates will be 1.6 times those listed, plus an additional line rental for the circuit between telephone exchanges which varies with the distance involved.

Table II

Communications Interfaces

DL11	Asynchronous single line interface. Speed up to 1200 bits per second with modems, up to 9600 b/s without modems. Full networks protocol must be implemented in software in the PDP-11.	$700
DUP11	Synchronous single line interface. Medium speed up to 9600 b/s. Performs CRC-16 block check calculations.	$1300
DMC11	Synchronous single line interface. High speed up to 19,200 b/s. Implements complete DDCMP protocol in hardware.	$2800

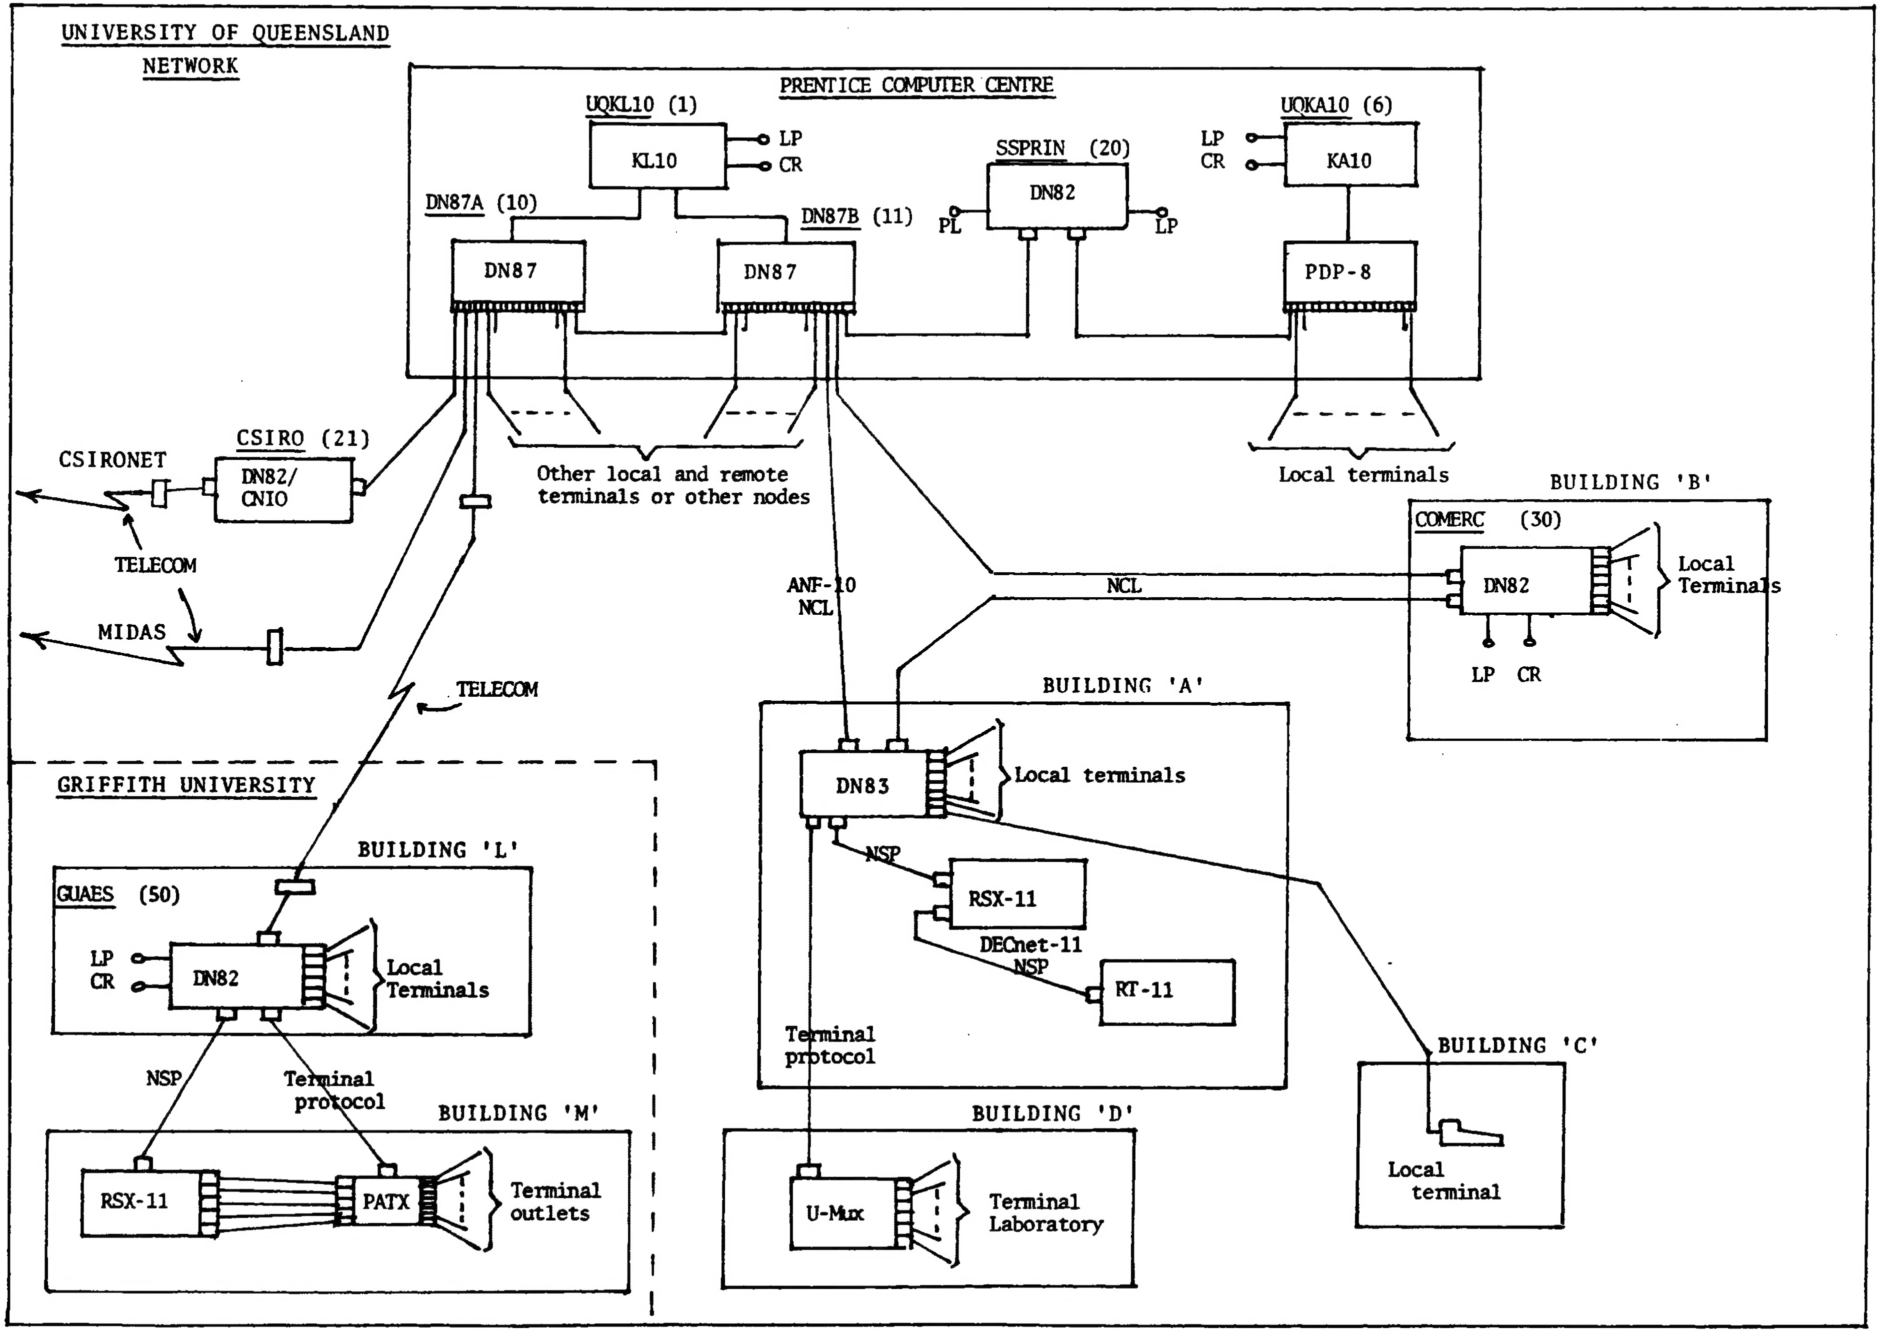
UNIVERSITY OF QUEENSLAND NETWORK
PRENTICE COMPUTER CENTRE
UQKL10 (1)
KL10
LP
CR
SSPRIN (20)
DN82
PL
LP
UQKA10 (6)
LP
CR
KA10
DN87A (10)
DN87
DN87B (11)
DN87
PDP-8
CSIRO (21)
DN82/
CNIO
CSIRONET
TELECOM
MIDAS
Other local and remote terminals or other nodes
Local terminals
BUILDING 'B'
COMERC (30)
DN82
Local Terminals
LP
CR
NCL
ANF-10
NCL
TELECOM
BUILDING 'A'
DN83
Local terminals
NSP
RSX-11
DECnet-11
NSP
RT-11
Terminal
protocol
GRIFFITH UNIVERSITY
BUILDING 'L'
GUAES (50)
LP
CR
DN82
Local
Terminals
NSP
Terminal
protocol
BUILDING 'M'
RSX-11
PATX
Terminal
outlets
BUILDING 'D'
U-Mux
Terminal
Laboratory
BUILDING 'C'
Local
terminal